A MUTUAL REINSURANCE SCHEME*

HARALD BOHMAN
Docent, University of Stockholm

ULF GRENANDER
L. Herbert Ballou University
Professor, Brown University

1. INTRODUCTION

1.1. *What is insurance?*

In order to make this report clear to those without experience in insurance matters, we first present some basic facts about the insurance business. In so doing we intentionally omit certain facts irrelevant to the present study. The most important omission of this kind is our assumption that the insurance business operates without administrative expenses and without sales costs. We also assume that the insurance business is run in such a way that no profit is made. It will be evident to readers already directly connected with insurance problems what further omissions we have made and why we have made them.

Insurance is the establishment of a contract between the insurance company and the insured person. This contract is by tradition called the "insurance policy" and the insured person is called the "policyholder". By agreeing to the insurance policy, the policyholder commits himself to paying certain premiums to the insurance company, and the insurance company commits itself to paying certain amounts to the policyholders. The conditions under which such payments are to be made are of many different kinds: the policyholder dies, the policyholder becomes ill, a homeowner's house is burnt down, or the policyholder has a collision in his car. The circumstances under which a payment is to be made to the policyholder are described in detail in the insurance policy. The amount to be paid is either fixed or variable: in life insurance the amount to be paid is always fixed and stated in the insurance policy, but in most

*) This is an abbreviated version. The full text is obtainable from Division of Applied Mathematics, Brown University, Providence, R.I. 02912, U.S.A.

other forms of insurance the intention is that the insurance compensate the policyholder for the losses he might incur as a consequence of the events covered by his policy.

How the premiums for individual policies should be established is an interesting problem in itself, but one which we must ignore here. In the present context we shall simply suppose the total premiums paid to the insurance company during one year to be equal to the expected claim payments. We further assume that the premiums make up the total income of the company, and that the claim payments make up all its expenses. At the beginning of each year, the insurance company can predict fairly well the amount of premiums that will be paid in the coming year, knowing the trend from earlier years and also how much they will spend for advertising, increased sales effort, etc.; thus, we can treat the payment of premiums as a deterministic process in this context. The stochastic part of the business is then considered to be the claim payments.

At the start of each year, then, the insurance company can foresee a certain amount of premium income during the coming year, and this amount will be equal to the expected amount of the claims. In the long run the premiums and the claims will thus be equal. But variations between individual years must be provided for. Hence, an insurance company must have at its disposal an equalization fund. Years during which the premium income is larger than the claim payments increase the equalization fund; years during which claims are larger than premiums will decrease the equalization fund. It is essential for the insurance company that the equalization fund never become negative, since this would mean that the company could not fulfill its obligation to pay the claims. This situation is traditionally called "ruin" in insurance literature. In stochastic terms, it is essential for the company that its "probability of ruin" be low.

There are many reasons why the equalization fund may decrease —e.g., the premiums may be set too low in relation to the claims. The reason that concerns us here, however, is the fluctuation of the claim payments. In order to reduce such fluctuations, so-called "reinsurance" is used. In re-insurance an insurance company agrees with other companies to take over certain portions of individual risks. The companies which do so are called the "reinsurers".

If a claim occurs on a reinsured risk, the reinsurer will pay its portion of the corresponding claim amount. In order that the reinsurer be willing to commit himself to such payments, he must receive a certain proportion of the corresponding premiums paid to the direct insurer. These portions of the premiums are usually called the "reinsurance premiums".

The reason for reinsurance may be formulated stochastically as follows. The insurance company enters a reinsurance contract in order to reduce the variance of the claim payments. *In this study we will suggest an alternative to the classical form of reinsurance, the efficiency of which will be measured by the reduction it produces in this variance.*

The general idea of this scheme is the following. A group of insurance companies agrees to create among themselves a mutual scheme called a Pool: in joining this Pool, the participating companies commit themselves to dividing certain claim amounts between them. For each participant company a limit GL is fixed. The company must report to the Pool all claims exceeding GL and occurring during one and the same year. These excess claim amounts are paid to the company in question, and these payments will be called "claim payments from the Pool". The sum of all amounts so reported will be divided between the participating companies according to certain rules to be described below. The amounts paid by the participating companies to the Pool in accordance with these rules will be called "premium payments to the Pool". If the rules for the division of the total claim amounts between participating companies are suitably devised, the whole scheme will result in a reduction of the variance of the claim payments.

The administration of the Pool will be simple. The result of the whole scheme is that the participating companies have to pay certain amounts between them each year. There will be no accumulation of funds. All that is needed is an organization which can undertake the necessary calculations on the data delivered by participating companies according to the regulations of the scheme.

An important restriction we have put upon ourselves in this study is pointed out here. It is often said that reinsurance is a levelling out of the risk result "in two dimensions". The Pool we shall discuss will have this levelling effect in one dimension only. A given company having extremely large claims in a given year might hope that not all of the other participating companies have the same high level of claim payments. During such a year, then, the other companies can take over a certain proportion of the first company's payments, and the levelling of the risk result is achieved. Suppose, however, that during a certain year most companies have such extremely large claim payments that even after the total claim payments have been distributed the net result is still a fairly high level of claim payments for each participating company. In order to overcome this difficulty, the Pool might establish an equalization fund, thus achieving the levelling out of the risk result between years usually considered to be one of the tasks of traditional reinsurance. If the Pool establishes an equalization fund, however, certain problems of a rather intricate nature must be solved. Perhaps the most difficult is to set up suitable rules concerning how large a proportion of the fund a company will be allowed to take with it if it wants to leave the Pool. A corresponding problem arises, of course, if a new member wants to enter the Pool. If the Pool is allowed to establish an equalization fund, this will no doubt increase its ability to level out the risk result of the participating companies. These problems seem worthy of a special study, but we have deliberately restricted ourselves in this study to a Pool with no equalization fund.

1.2. *Methods used in this study*

Because of the exploratory nature of our study it seemed desirable to do the computing interactively, and we decided to do all of it in APL. This very powerful programming language is a great help in using the computer to test ideas during the development of a mathematical model. The present problem is an excellent example of one which may be solved by using the computer as an experimental laboratory.

In the present case we knew from the very beginning certain criteria which must be met by the model. In order that insurance companies be willing to participate in the scheme, for example, it would be necessary that the expected value of payments to the Pool be equal to the expected value of payments from the Pool for each participating company; *the system should be fair on a net basis.*

In order that the companies find the scheme advantageous, *it is also necessary that for each participating company the variance of the payments from the Pool be larger than the variance of the payments to the Pool.* It could also be foreseen that no scheme could be agreed upon were not an upper limit applied for the amounts which each participating company is allowed to report to the Pool. These examples may serve to illustrate the fact that many factors had to be taken into consideration and that it would have been impossible to predict the effect of the different criteria using classical analytical methods. It is at this point that the computer comes in. When a first rough model has been worked out, its results may be tested on the computer; the model may then be modified and tested again, as often as desired. This is a problem-solving method which was not available in the precomputer era. It goes without saying that the APL language, with its sophisticated sets of instructions and its highly efficient system for changing already-written programs, is a powerful tool in such a study.

2. Mathematical Models

2.1. *Eight insurance companies and their portfolios*

While we were working on the scheme we tested the various ideas by simulating them on the computer. It was therefore necessary for us to have data from a group of insurance companies assumed to be members of the Pool. The model used for this purpose will now be described.

We assume that eight companies are members of the Pool. Each company classifies its claims into small claims, medium-sized claims and large claims. Within each class the amounts of the claims have an exponential distribution, the parameter of which is chosen in such a way that the mean value of small claims is 1, the mean value of medium-sized claims is 10 and the mean value of large claims is 100.

Table 2.1.1

	Expected number of claims							
ompany no	1	2	3	4	5	6	7	8
mall	1000	1000	1000	1000	10,000	10,000	10,000	10,000
edium	100	100	500	500	1,000	1,000	5,000	5,000
arge	10	50	50	250	100	500	500	2,500
um	1100	1150	1550	1750	11,100	11,500	15,500	17,500

It is thus envisioned that the distribution function corresponding to claims within the same class will be the same in all companies.

The eight companies, which will be designated below by the numbers 1-8, vary in size in different respects. The expected number of claims in company 5 is ten times the expected number of claims in company 1. The same relation holds for companies 6 and 2, and so on. The companies also vary among themselves in the proportion of medium- and large-sized claims to the total number of claims; this may be seen in Table 2.1.1, where the expected number of claims in each category is given for each company.

The actual claim results of the different companies can be looked upon as the outcome of a stochastic process in which a number of random mechanisms are at work. One such mechanism is that which gives the actual number of claims during a certain year. It is assumed that for each company and each category of claims the number of claims follows a Poisson distribution the parameter of which is given in Table 2.1.1. It is also assumed that all claims of all categories occur independently of each other.

The second random mechanism serves to determine the size of a claim. This mechanism gives the result of a random experiment where the variable has the distribution function $(1 - \exp(-x/M))$ with M equal to the mean value of claims in the category in question.

From this model we can now calculate, for each company and each year, the actual number of claims within each category and also the size of each individual claim.

We have described the model of the eight companies here in exactly the same way as we applied the random mechanisms while performing the simulation. There is another interpretation of the same model, stochastically equivalent but more in line with the traditional method of presentation, which we shall describe briefly. One random mechanism generates the actual number of claims, irrespective of the size of the claim This mechanism gives for company 2, for example, the actual number of claims, knowing that the expected number is equal to 1150. The second random mechanism gives the size of each individual claim. In doing this it operates on the distribution of all claims, regardless of whether they are large, medium-sized or small. This distribution function is clearly

TABLE 2.1.2

x/MEAN	Companies 1 and 5	Companies 2 and 6	Companies 3 and 7	Companies 4 and 8
	$1-F$			
1	.1379	.0902	.1892	.1683
2	.0650	.0642	.1060	.1085
4	.0387	.0417	.0432	.0706
6	.0255	.0324	.0256	.0494
8	.0176	.0274	.0194	.0346
10	.0129	.0238	.0161	.0243
	Moments of claim distribution			
1st	2.70	6.09	7.10	17.71
2nd	200	889	711	2915
3rd	55,000	261,000	196,000	859,000

equal in our case to the weighted mean of the distribution functions for each category, and the weighting is done in proportion to the expected number of claims within each category. Some data concerning the distribution function obtained in this way are given in Table 2.1.2.

We have simulated the process for two different cases. In the first case the process is stationary, which means, *inter alia*, that the portfolios are exactly the same each year. In the second case we have assumed that there is a 5% increase in prices each year caused by inflation, which means that the mean values of the claims 1, 10 and 100 are increased by 5% each year the simulation is performed. At the same time, we have assumed that the structure of the portfolio changes so that the expected number of large and medium-sized claims also increases by 5% each year; this change is intended to reflect the possibility that the companies participating in the Pool may be willing to take on larger risks than earlier.

2.2. *The Pool*

For each participating company there is prescribed one lower limit, GL, and one upper limit, GU (see Table 2.2.1). If the company has a claim of size X, and X is larger than GL, then the company will receive the amount (X-GL) from the Pool. If, however, X is larger than GU, the upper limit applies and the Pool will pay only the amount (GU-GL). The total amount of claims placed in the Pool during one year is determined according to this rule.

TABLE 2.2.1

Company no.	1	2	3	4	5	6	7	8
GL	69	230	230	391	300	460	460	620
GU	690	2300	2300	3916	3000	4600	4600	6200

In the case in which the claim amounts increase because of inflation, GU and GL increase at the same rate.

Two criteria must be fulfilled in determining the portion of the Pool to be paid by each participating company:

The expected value of the premium payments P_i to the Pool should be equal to the expected value of the claim payments S_i to the company. This should hold for all companies.
The variance of P_i should if possible be less than the variance of S_i. This should also hold for all companies.

In order to fulfill these requirements we suggest that the following premium formula be applied:

$$P_i = m_i + Q \cdot \sigma_i \tag{2.2.1}$$

where m_i = Expected value $\{P_i\}$ and σ_i^2 = Variance $\{P_i\}$. This choice is discussed in section 3. The quantity Q is common to all companies and is determined by the following formula:

$$\Sigma S_i = \Sigma m_i + Q \Sigma \sigma_i \tag{2.2.2}$$

In order for the Pool to be able to calculate m_i and σ_i, it must have at its disposal the corresponding probability distribution. This is established by having each participating company inform the Pool of the size of its hundred largest claims for the past year. The Pool then fits a distribution function of the following three-parameter family to the data provided by the company:

$$1 - F(x) = p \cdot \exp\left\{\frac{t_1 - x}{a}\right\} + (1 - p) \exp\left\{\frac{t_1 - x}{b}\right\} \tag{2.2.3}$$

where

$$a = \frac{p \cdot q}{1 - p}$$

$$b = \frac{(1 - p)\, q}{p}$$

$$t_1 + q = \text{mean of distribution}$$

$$t_1 = \text{smallest claim among the 100 reported}$$

$$\text{Variance} = V^2 = 2\left\{\frac{(1 - p)^2}{p} + \frac{p^2}{1 - p}\right\} q^2 - q^2$$

$$p = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{2q^2}{V^2 + 7 \cdot q^2}} \tag{2.2.4}$$

This is only possible if $V^2 \geq q^2$. As a consequence we have put $p = 1/2$ when the hundred reported claims show a $V^2 < q^2$. This limitation seems to be of no practical importance.

When the parameters are chosen in this manner, the three-parameter distribution function will have the same mean value and variance as the hundred actual claims.

The value of m_i and σ_i now may be found from the following formulas, where for the sake of clarity we have omitted the index "i" indicating that the calculations are to be performed for each company individually.

$$m = 100 \int_{GL}^{GU} (X - GU)\, dF + 100(GU - GL)\, (1 - F(GU)) \tag{2.2.5}$$

$$m_2 = 100 \int_{GL}^{GU} (X - GU)^2\, dF + 100(GU - GL)^2\, (1 - F(GU)) \tag{2.2.6}$$

$$\sigma = \sqrt{m_2 - \frac{m^2}{100}} \tag{2.2.7}$$

The values of m_i and σ_i are calculated on the basis of the hundred claims presented by each company each year according to the method just described. It seems natural to ask if the variance of the estimates could be made still smaller by using estimates also from earlier years. We have therefore tried a method of estimating m_i and σ_i by weighting together the m_i and σ_i for the previous year and earlier years. The "weighted" values of m_i and σ_i are denoted by $\bar{m}_i$ and $\bar{\sigma}_i$.

$$\bar{m}_{i,\,k} = \frac{m_{i,\,k} + R\,.\,m_{i,\,k-1} + R^2 m_{i,\,k-1} + \ldots}{1 + R + R^2 + \ldots} \tag{2.2.8}$$

$$\bar{\sigma}_{i,\,k} = \frac{\sigma_{i,\,k} + R\,.\,\sigma_{i,\,k-1} + R^2\,.\,\sigma_{i,\,k-2} + \ldots}{1 + R + R^2 + \ldots} \tag{2.2.9}$$

where the summation is extended over the preceding years of the plan, and where $m_{i,k}$ - the m-value for company i in year k, $\sigma_{i,k}$ analogously. We have chosen the value $5\sqrt{\frac{1}{2}}$ for R, a halflife of five years. Note that this *adaptive method of estimation* can be expected to detect secular changes faster than for the usual estimate with $R = 1$.

2.3. *Why the Pareto distribution was not used*

It is often said that the Pareto distribution should be used as a claim distribution for the insurance business, since other distributions tend to give too optimistic a picture of the claims. We have discussed this question and *do not accept this widespread opinion*. As has already been said, we use certain exponential polynomials as distribution functions both when the Pool must fit the observed data to a suitable distribution and when simulating the claims of the eight companies. We will present our reasons for doing so by discussing the arguments at some length.

The large claims often make up a fairly large portion of the total claim costs in an insurance portfolio. The number of large claims, also, is often very small as compared to the number of small claims. Under such circumstances it is evident that estimates of the moments of the claim distribution are completely unreliable as soon as we go beyond the second moment. This means, in practice, that actual data must often be fitted to a class of distribution functions on the basis of the first two moments only.

Suppose we now proceed as follows. We start with the simple assumption that all claims are equal to 1. This means that the mean value is 1 and the variance is 0. Suppose now, for example, that the observed values are mean value = 1, variance 25. We modify the first choice of distribution function by adding, with probability .001, a claim of size 160, while all the other claims are changed to be of size .84 instead of 1. This is a two-point distribution with probability .999 that $X = .84$ and probability .001 that $X = 160$. This distribution no doubt has approximately mean value = 1, variance = 25, but nobody would recommend that the procedure described here be used to fit the actual data to a suitable distribution function. This becomes even more relevant if we also suppose that we have claims from another company with mean value = 1 and variance = 50 and that we choose a distribution function equal to the one just discussed, changing 160 to 220 and .84 to .78.

TABLE 2.3.1

	$1 - F(10)$	
s.d.	Exponentials	Pareto
1	.000	.001
2	.008	.002
3	.022	.002
4	.030	.002
5	.030	.002
6	.028	.002
7	.024	.002
	$1 - F(5)$	
s.d.	Exponentials	Pareto
1	.005	.006
2	.038	.009
3	.053	.009
4	.051	.010
5	.044	.010
6	.036	.010
7	.029	.010

It could be said that the procedure described above is a caricature of how curve-fitting should be done. There is, however, much in it which applies to the fitting of a Pareto distribution. To demonstrate this we calculated the values of the distribution function according to Pareto, choosing the parameters so that the mean value = 1 and the standard deviation = 1, 2, . . ., 7. The results are given in Table 2.3.1, as are the corresponding figures for a distribution function of the type in Equation (2.2.3) with the same mean value and standard deviation. It may be seen from this table that the Pareto distributions for higher values of X are approximately the same for all values of the standard deviation given. These figures suggest, to state it a bit imprecisely, that the Pareto distribution is changed to higher values of the variance not by making the "tail" thicker or thinner but by making the very unexpected large claim still larger. The figures look the way they do because of the formula for the standard deviation of the Pareto distribution:

$$\begin{cases} 1 - F(x) = \left(\frac{\alpha}{x}\right)^{\beta} \quad \text{for } x > \alpha. \\ \frac{\text{s.d.}}{\text{mean}} = \frac{1}{\sqrt{\beta(\beta - 2)}} \end{cases}$$

In order to get large values of the standard deviation we must choose β close to 2, which means that the tail of the distribution will always look approximately the same as soon as the standard deviation is large enough. This is exactly the same situation as in our first two simplified cases.

3. Determination of the Premiums

We now turn to the problem of how to find the "best" values for the premiums P_i. The following formula has been used

$$P_i = m_i + \sigma_i \frac{\Sigma(S_j - m_j)}{\Sigma\sigma_j}$$

In what sense this premium can be said to be the "best" one is not explaned in this abbreviated version The full discussion is found in the original version.

4. How the Simulations Were Carried Out

One straightforward way of performing the simulations would be to proceed in the following manner For each company, and for each category of claim, simulate the number of claims for a given year according to the Poisson distribution in question. Then for each category perform a simulation as many times as the actual number of claims indicates to get the corresponding claim amounts.

Having thus arrived at the total number of actual claims and all individual claim amounts, order the claims with respect to size. Then select the hundred largest claims and report them to the Pool. The rest of the work is, of course, purely routine. However, it was evident that this straightforward method would have taken considerable computer time, and we thus had to choose a better solution.

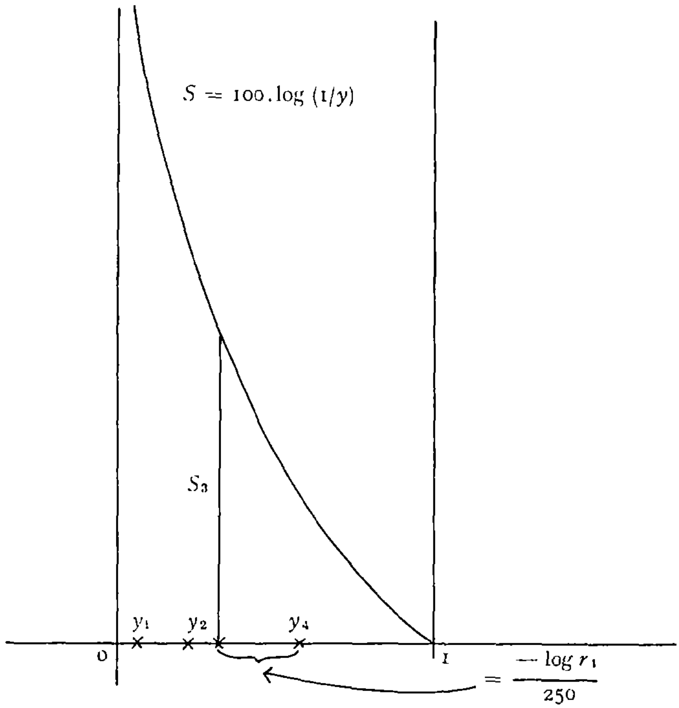

To illustrate the method actually used in the simulation, let us choose the large claims of company 4. The expected number of these claims is 250 The procedure may be more easily understood by reference to Figure 4.1. We start at point 0 and proceed stepwise towards point 1. Consider the way in which customers are assumed to arrive in queueing theory: they arrive independently of each other and the expected number of customers that actually arrive in one hour is Poisson-distributed. Each step is equal to the time between arrival of one customer and the next, and the length of the step is obtained by taking the negative logarithm of a random number evenly distributed in (0,1) and then dividing by 250, the number of customers (or large claims). The number of points then obtained in (0,1) will be equal to the actual number of the large claims.

The curve $100 . \log (1/y)$ in Figure 4.1 is of course the inverse of the corresponding claim distribution. It is evident that if we take, for example, the distance from point 3 in Figure 4.1 up to this curve, it will be equal to the third largest claim obtained in this category. If we use this method, we can clearly stop after having simulated 100 steps from point 0, since each such step corresponds to one claim and since only the hundred largest claims need be reported to the Pool.

When we carried out this procedure we actually applied the same idea to all three categories of claims at once, each time deciding, on the basis of the claim last chosen, if the next step should be "small claim step", "medium-sized claim step" or "large claim step". By so doing we could pick out the hundred largest claims merely by choosing 100 random numbers and calculating the corresponding amounts.

The process was simulated for a period of 25 years; after completion, the observed values of the mean and variance of P_t and S_t were calculated.

5. Conclusions

The results for the individual years are given in the Appendix. One simulation corresponds to the assumption of stationarity, and one corresponds to the 5% yearly increase of certain factors, as described earlier. For each simulation P is calculated both weighted and not weighted, where "weighted" has the meaning used in equations (2.2.8) and 2.2.9). The results for the 25 years are given in Table 5.1 below. It may be seen that the criterion that the expected values of P_t and S_t should be equal is fairly well fulfilled. It also may be seen that the second criterion, that the variance of P_t should if possible be less than that of S_t, is also reasonably well fulfilled, and that is especially true in the case in which the weighted premiums were used.

Following the 25 years described in the data given in the Appendix gives a more detailed picture than does Table 5.1 of the economic functioning of the mutual Pool, and the reader is encouraged to scrutinize the data for the individual years.

TABLE 5.1

Company no.	1	2	3	4	5	6	7	8
	The stationary case: results after 25 years							
Mean value of								
Claims	614.7	374.3	526.8	579.4	420.2	471.3	514.1	447.8
Premiums weighted	517	355.4	461.2	630.9	460.9	493.9	557.7	471.9
Premiums unweighted	585.6	364.7	489.3	584	454.1	470	520.6	480.3
Variance of								
Claims	1.1E5	7.5E4	8.0E4	1.7E5	9.8E4	4.7E4	6.4E4	7.5E4
Premiums weighted	3.1E4	1.1E4	2.0E4	2.9E4	1.8E4	1.4E4	1.6E4	1.6E4
Premiums unweighted	1.3E5	4.5E4	6.2E4	1.2E5	7.2E4	2.0E4	4.2E4	3.8E4
	The case with a trend: results after 25 years							
Mean value of								
Claims	1893	2188	1702	2049	1859	2015	2379	2062
Premiums weighted	1766	2169	1836	2049	2092	2036	2218	1980
Premiums unweighted	1731	2103	1687	2056	1909	2146	2384	2040
Variance of								
Claims	1.8E6	3.3E6	1.6E6	2.3E6	1.7E6	3.3E6	4.8E6	3.2E6
Premiums weighted	1.4E6	2.2E6	1.2E6	1.9E6	1.7E6	2.1E6	3.2E6	2.1E6
Premiums unweighted	1.5E6	2.7E6	1.4E6	2.4E6	2.0E6	3.2E6	4.2E6	2.9E6